

CLAIMS

I/We claim:

- [c1] 1. A communications architecture for communicating between hosts and data store devices, the communications architecture comprising:
- a plurality of hosts, each host having a communications interface with a serial communications link;
 - a plurality of data store devices, each data store device having a communications interface with a serial communications link; and
 - a switching network having communications interfaces with serial communications links for establishing communications paths between hosts and data store devices
- wherein the communications architecture supports control packets and data packets, the control packets and data packets having headers with different formats, supports data packets having a variable length, and supports preemption of data packets by control packets.
- [c2] 2. The communications architecture of claim 1 wherein the hosts, data store devices, and switching network form a storage area network.
- [c3] 3. The communications architecture of claim 1 wherein the data store devices are disk-based.
- [c4] 4. The communications architecture of claim 1 wherein the data store devices are block-oriented.
- [c5] 5. The communications architecture of claim 1 wherein the communications architecture supports dynamic segmentation of data packets.

- [c6] 6. The communications architecture of claim 5 wherein the communications architecture supports merging of segmented data packets.

- [c7] 7. The communications architecture of claim 1 wherein multiple data packets form a transaction, the data packets having an order within the transaction, and wherein the communications architecture ensures that the data store devices receive data packets of a transaction in order.

- [c8] 8. The communications architecture of claim 7 wherein the communications architecture does not ensure that the hosts receive data packets in order of transmission.

- [c9] 9. The communications architecture of claim 1 wherein when an error is detected during processing of a transaction, the host that initiated the transaction is notified so that the initiating host can handle the error.

- [c10] 10. The communications architecture of claim 9 wherein when an error is detected by the switching network, the switching network forwards an error message to the initiating host.

- [c11] 11. The communications architecture of claim 9 wherein the error is at a physical layer.

- [c12] 12. The communications architecture of claim 9 wherein the error is at a link layer.

- [c13] 13. The communications architecture of claim 9 wherein the error is at a transport layer.

[c15] 15. The communications architecture of claim 14 wherein the switching network does not attempt to correct errors.

[c16] 16. The communications architecture of claim 1 wherein the switching network asymmetrically processes packets received from the hosts and packets received from the data store devices.

[c17] 17. The communications architecture of claim 16 wherein the asymmetric processing includes handling of errors.

[c18] 18. The communications architecture of claim 16 wherein the asymmetric processing includes ordering of packets within a transaction.

[c19] 19. A communications architecture for communicating between hosts and data store devices, the communications architecture comprising:

a plurality of hosts, each host having a communications interface with a serial communications link;

a plurality of data store devices, each data store device having a communications interface with a serial communications link and being block-oriented; and

a switching network having communications interfaces with serial communications links for establishing communications paths between hosts and data store devices, the switching network for asymmetrically processing packets transmitted by hosts and packets transmitted by data store devices.

[c20] 20. The communications architecture of claim 19 wherein the asymmetric processing of packets includes ensuring that packets of a transaction transmitted by a host are received in order by a data store device and not ensuring that packets of a transaction transmitted by a data store device are received in order by a host.

[c21] 21. The communications architecture of claim 20 wherein packets of a transaction are guaranteed to be transmitted along the same path from a host to a data store device and not guaranteed to be transmitted along the same path from a data store device to a host.

[c22] 22. The communications architecture of claim 20 wherein packet-level load balancing is performed on packets transmitted from a data store device to a host, but not performed on packets transmitted from a host to a data store device.

[c23] 23. The communications architecture of claim 19 wherein the asymmetric processing of packets includes when an error is detected during transmission of a packet from a host to a data store device notifying the host of the error so that the host can handle the error and when an error is detected during transmission of a packet from a data store device to a host notifying the host of the error so that the host can handle the error.

[c24] 24. The communications architecture of claim 23 wherein only the host that receives the notification performs error handling.

[c25] 25. The communications architecture of claim 19 wherein the hosts, the data store devices, and the switching network form a storage area network.

[c26] 26. The communications architecture of claim 19 wherein the data store devices are disk-based.

[c27] 27. The communications architecture of claim 19 wherein the communications architecture supports control packets and data packets, the control packets and data packets having headers with different formats.

[c28] 28. The communications architecture of claim 19 wherein the communications architecture supports data packets having a variable number of blocks.

[c29] 29. The communications architecture of claim 19 having control packets and data packets and wherein transmission of data packets can be preempted by control packets.

[c30] 30. A network for transmitting data between hosts and data store devices, the data store devices being block-oriented, the switching network comprising:

- a plurality of communications interfaces to hosts;
- a plurality of communications interfaces to data store devices; and
- a path component for establishing communications paths between hosts and data store devices, for transmitting data packets and control packets, the control packets and data packets having headers with different formats, the data packet having a variable length, and for preempting transmission of data packets to transmit control packets.

[c31] 31. The network of claim 30 wherein the hosts, data store devices, and network form a storage area network.

[c32] 32. The network of claim 30 wherein the path component dynamically segments data packets.

[c33] 33. The network of claim 32 wherein the path component dynamically merges segmented data packets.

[c34] 34. The network of claim 32 wherein multiple data packets form a transaction, the data packets having an order within the transaction, and wherein the path component ensures that the data store devices receive data packets of a transaction in order.

[c35] 35. The network of claim 34 wherein the path component does not ensure that the hosts receive data packets of a transaction in order.

[c36] 36. The network of claim 30 wherein when an error is detected by the path component during processing of a transaction, the host that initiated the transaction is notified so that the initiating host can handle the error.

[c37] 37. The network of claim 36 wherein the path component does not attempt to correct errors.

[c38] 38. The network of claim 30 wherein the path component asymmetrically processes packets received from the hosts and packets received from the data store devices.

[c39] 39. The network of claim 38 wherein the asymmetric processing includes handling of errors.

[c40] 40. The network of claim 38 wherein the asymmetric processing includes ordering of packets within a transaction.